

Subcubic edge chromatic critical graphs have many edges

Daniel W. Cranston* Landon Rabern†

June 16, 2015

Abstract

We consider graphs G with $\Delta = 3$ such that $\chi'(G) = 4$ and $\chi'(G - e) = 3$ for every edge e , so-called *critical* graphs. Jakobsen noted that the Petersen graph with a vertex deleted, P^* , is such a graph and has average degree only $2 + \frac{2}{3}$. He showed that every critical graph has average degree at least $2 + \frac{2}{3}$, and asked if P^* is the only graph where equality holds. We answer his question affirmatively. Our main result is that every subcubic critical graph, other than P^* , has average degree at least $2 + \frac{26}{37} = 2.\overline{702}$.

1 Introduction

A *proper edge-coloring* of a graph G assigns a color to each edge in $E(G)$ such that edges with a common endpoint receive distinct colors. The minimum number of colors needed for a proper edge-coloring is the *edge-chromatic number* of G , denoted $\chi'(G)$. The maximum degree of G is denoted $\Delta(G)$, or simply Δ when G is clear from context. Note that always $\chi'(G) \geq \Delta(G)$. Vizing famously proved that $\Delta(G) + 1 \geq \chi'(G) \geq \Delta(G)$ for every graph G . A graph is *edge-chromatic critical* (also Δ -*critical*, or simply *critical*) if $\chi'(G) > \Delta(G)$ but $\chi'(G - e) = \Delta(G)$ for every edge e . A vertex of degree k is a k -*vertex*. If v_1 is a k -vertex and v_1 is adjacent to v_2 , then v_1 is a k -*neighbor* of v_2 .

Vizing [8, 9] was the first to seek a lower bound on the number of edges in a critical graph, in terms of its number of vertices. This problem is now widely studied, for a large range of maximum degrees Δ . Woodall gives a nice history of this work, in the introduction to [10]. In this paper, we study the problem for subcubic graphs, i.e., when $\Delta = 3$.

It is easy to check that 2-critical graphs are precisely odd cycles, which are completely understood. So the first non-trivial case is 3-critical graphs. Let P^* denote the Petersen graph with a vertex deleted. Jakobsen [4] observed that P^* is 3-critical and has average degree $2 + \frac{2}{3}$. In the same paper, he asked if every 3-critical graph has average degree at least $2 + \frac{2}{3}$. A year later [5], he answered this question affirmatively. However, in this second paper Jakobsen asked whether this bound holds with equality for any graph other than P^* .

*Department of Mathematics and Applied Mathematics, Virginia Commonwealth University, Richmond, VA; dcranston@vcu.edu; Research of the first author is partially supported by NSA Grant 98230-15-1-0013.

†LBD Data Solutions, Lancaster, PA; landon.rabern@gmail.com




Figure 1: P^* and Woodall's first two examples: J_1 , and J_2 ; the first three examples in an infinite family J_i of 3-critical graphs with $2|E(G)| < (2 + \frac{3}{4})|V(G)|$.

A natural extension of this question is to determine the minimum α such that there exists an infinite sequence of 3-critical graphs with average degree at most $2 + \alpha$. The first progress toward answering this question is due to Fiorini and Wilson [3, p. 43], who constructed an infinite family of 3-critical graphs with average degree approaching $2 + \frac{3}{4}$ from below. Woodall [10, p. 815] gave another family with the same number of edges and vertices; see Figure 1. Before presenting his construction, we need a definition.

Let G_1 and G_2 be two graphs with $v_1v_2 \in E(G_1)$ and $v_3v_4 \in E(G_2)$. A *Hajós join* of G_1 and G_2 is formed from the disjoint union of $G_1 - v_1v_2$ and $G_2 - v_3v_4$ by identifying vertices v_1 and v_3 and adding the edge v_2v_4 .

Lemma 1. *If G_1 and G_2 are k -critical graphs, and G is a Hajós join of G_1 and G_2 that has maximum degree k , then G is also k -critical.*

This is an old result of Jakobsen [4]. It is a straightforward exercise, so we omit the details, which are available in Fiorini & Wilson [2, p. 82–83] and Stiebitz et al. [7, p. 94].

Corollary 2. *Let G_1 and G_2 be subcubic graphs, and let G_1 be 3-critical. If G is a subcubic graph that is a Hajós join of G_1 and G_2 , then G is 3-critical if and only if G_2 is 3-critical.*

Proof. The “if” direction follows immediately from the previous lemma. To prove the “only if” direction, we can assume that $\chi'(G_2) = 3$ and construct a 3-coloring of G from 3-colorings of G_2 and $G_1 - v_1v_2$. \square

Now we present Woodall's construction of 3-critical graphs.

Example 1. *Form J_k by starting with P^* and taking the Hajós join with P^* a total of k times (successively), so that each intermediate graph has $\Delta = 3$. The resulting graph J_k is 3-critical, has $11k + 12$ edges and $8k + 9$ vertices. Thus, the average degree of the sequence J_k approaches $2 + \frac{3}{4}$ from below.*

The vertex and edge counts follow immediately by induction. That J_k is 3-critical uses induction and also Lemma 1. \square

Our main result is that every 3-critical graph G , other than P^* , has average degree at least $2 + \frac{26}{37}$. Before we prove this, it is helpful to provide a brief outline. Our proof uses the discharging method. More precisely, we first show that numerous subgraphs, not necessarily

induced, are forbidden from appearing in a minimal counterexample G . To conclude, we give each vertex a charge equal to its degree. Under the assumption that G contains none of the forbidden subgraphs, we redistribute charge, without changing its sum, so that each vertex has final charge at least $2 + \frac{26}{37}$. This proves the theorem.

The intuition behind our proof is to show that every vertex in the graph has many “nearby” 3-vertices and not too many nearby 2-vertices. Let H be the subgraph of G induced by 3-vertices with 2-neighbors. To facilitate the discharging, we show that each component of H is small. Further, for each 3-vertex v not in H , we show that the sum of the sizes of the adjacent components of H is small, so v can give much of its extra charge to each vertex in these components of H .

We mentioned above that our proof begins by forbidding certain subgraphs from appearing in a critical graph. The easiest example of this is that no critical graph contains a 1-vertex. If so, we delete its incident edge e , color $G - e$ by criticality, then extend the coloring to e . More standard examples often require that we recolor part of the graph before we extend the coloring. An (x, y) -Kempe chain is a component of the subgraph induced by the edges colored x and y . Note that each Kempe chain is either a path or an even cycle. If vertices v_1 and v_2 lie in the same (x, y) -Kempe chain, then v_1 and v_2 are x, y -linked. Given a proper coloring of (some subgraph of) a graph G , if we interchange the colors on some (x, y) -Kempe chain, the resulting coloring is again proper. This interchange is an (x, y) -Kempe swap and plays a central role in most proofs of forbidden subgraphs in critical graphs. If a color w is used on an edge incident to a vertex v , then we say that v sees w ; otherwise v misses w .

Suppose that $d(v_1) = 2$, $d(v_2) = 3$, and $v_1v_2 \in E(G)$. Suppose also that we 3-color $G - v_1v_2$ with colors x , y , and z . If we cannot extend this coloring to G , then (by symmetry) we may assume that v_1 sees x and that v_2 sees y and z . Furthermore, v_1 and v_2 must be x, y -linked; otherwise we perform an (x, y) -Kempe swap at v_1 and afterwards color v_1v_2 with x . Similarly, v_1 and v_2 must be x, z -linked. The quintessential tool for forbidding a subgraph in a critical graph is Vizing’s Adjacency Lemma, which he proved using Kempe chains and a similar structure for recoloring, known as Vizing fans.

Vizing’s Adjacency Lemma. *Let G be a Δ -critical graph. If $u, v \in V(G)$ and $uv \in E(G)$, then the number of Δ -neighbors of u is at least $\max\{2, \Delta - d(v) + 1\}$.*

The proof is available in Fiorini & Wilson [2, p. 72–74] and in Stiebitz et al. [7]. In the case $\Delta = 3$, we conclude that every 3-vertex has at most one 2-neighbor. This is helpful in our goal to prove that every 3-vertex has many nearby 3-vertices.

Two of our proofs that certain subgraphs are forbidden from G are a bit lengthy. To improve readability, we simply state the results when we need them in Section 2 (as Claims 2 and 5), and defer the proofs to Section 3. By using a computer, we were able to improve our edge bound for 3-critical graphs to $2|E(G)| \geq (2 + \frac{42}{59})|V(G)|$. However, a human-readable proof is too long to include here (roughly 100 pages). We discuss this work a bit more in Section 4, as well as give a web link where that proof is available.

2 Proof of Main Theorem

Main Theorem. *Let P^* denote the Petersen graph with a vertex deleted. If a graph G with $\Delta = 3$ is critical, then either $2|E(G)| \geq (2 + \frac{26}{37})|V(G)|$ or else either G is P^* .*

Proof. We will prove the following variation, from which the version stated above follows: If G is 3-critical then either $2|E(G)| \geq (2 + \frac{26}{37})|V(G)|$ or else $G = P^*$ or G is the Hajós join of P^* and a smaller 3-critical graph. To see that this implies the version stated above, consider a smallest 3-critical graph G . Since it is smallest, either G has average degree at least $2 + \frac{26}{37}$ or else G is P^* . Note that if G has average degree less than $2 + \frac{3}{4}$, then the Hajós join of G and P^* has average degree higher than that of G (since it has 6 more 3-vertices and two more 2-vertices). Thus, every 3-critical graph with average degree less than $2 + \frac{26}{37}$ must be formed from P^* by repeatedly taking the Hajós join with further copies of P^* ; these are Woodall's construction, from Example 1. It is easy to check that already J_1 has average greater than $2 + \frac{26}{37}$. Now we prove this variation of the theorem.

Suppose the theorem is false, and let G be a minimal counterexample. Note, as follows, that G is 2-edge-connected, so has minimum degree 2. If G is disconnected, then we can color each component by minimality. Similarly, suppose G has a cut-edge v_1v_2 . By minimality, we can color $G - v_1v_2$, and permute the colors so that the same color is missing from v_1 and v_2 . Before giving the discharging, we prove some structural claims about G and H .

Claim 1. *G has no adjacent 2-vertices, no 3-vertex with two or more 2-neighbors, and no 3-cycle.*

If G has adjacent 2-vertices v_1 and v_2 , then color $G - v_1v_2$. Now at most two colors are forbidden on v_1v_2 , so we can extend the coloring. Recall that Vizing's Adjacency Lemma guarantees that each 3-vertex has at least two 3-neighbors, so at most one 2-neighbor. Now suppose that G has a 3-cycle $v_1v_2v_3$. First, suppose that v_1v_2 lies in a second 3-cycle $v_1v_2v_4$. If v_3v_4 is also in G , then $G \cong K_4$, so $\chi'(G) = 3$. So suppose not. Let v_5 be a neighbor of v_3 other than v_1 and v_2 (or a neighbor of v_4 other than v_1 and v_2). To form G' from G , contract $\{v_1, v_2, v_3, v_4, v_5\}$ to a single vertex. Since G has no cut-edges, each of vertices v_1, v_2, v_3, v_4 has degree 3; thus, the average degree of G' is less than that of G . Now by minimality, $\chi'(G') = 3$, and we can extend this coloring of G' to G . So we may assume that no edge of $v_1v_2v_3$ lies on a second triangle. To form G' from G , contract the three edges of triangle $v_1v_2v_3$. Again, by minimality, $\chi'(G') = 3$, and we can extend the coloring to G . \square




Figure 2: Three subgraphs forbidden from a 3-critical graph G . Vertices drawn as rectangles have degree 2 in G and those drawn as circles have degree 3 in G .

Claim 2. *The subgraphs shown in Figures 2(a), 2(b), and 2(c) are forbidden.*

The proof for Figure 2(a) is given in Lemma 4 in Section 3. For Figures 2(b) and 2(c), we give the proof here. We begin with Figure 2(b). By criticality, we use colors x , y , and z to color $G - v_3v_5$; call the coloring φ . WLOG, v_5 sees x , $\varphi(v_2v_3) = y$, and $\varphi(v_3v_4) = z$. If v_1 misses x , then $\varphi(v_1v_2) = z$ and $\varphi(v_1v_4) = y$. Now do an (x, y) -Kempe swap at v_5 . Edge v_3v_5 will be colorable unless the (x, y) -Kempe path starting at v_5 ends at v_3 , so assume that it does. Now v_5 sees y and v_1 sees y . Thus, in the original case, we can assume that v_1 sees x . WLOG, $\varphi(v_1v_2) = x$. Uncolor the edges incident to v_1 and v_3 and color v_2v_3 with x . Now greedily color v_2v_1 , v_1v_4 , v_4v_3 , and v_3v_5 .

Now consider Figure 2(c). As above, we use colors x , y , z to 3-color $G - v_3v_6$; call the coloring φ . WLOG, v_6 sees x . Since $\varphi(v_1v_2) \neq \varphi(v_1v_4)$, by symmetry, assume $\varphi(v_2v_5) = x$. Again, WLOG, $\varphi(v_2v_3) = y$, $\varphi(v_3v_4) = z$, and $\varphi(v_1v_2) = z$. We may assume that v_3 and v_6 are x, y -linked. Thus, v_5 sees y . If $\varphi(v_1v_4) = y$, then do a (y, z) -Kempe swap at v_3 (the entire component is just the 4-cycle $v_1v_2v_3v_4$). Now v_3 and v_6 are no longer x, z -linked, so do an (x, z) -Kempe swap at v_3 , and color v_3v_6 with z . Thus, we assume that $\varphi(v_1v_4) = x$. Now again, do an (x, z) -Kempe swap at v_3 , then color v_3v_6 with z . \square

Recall that H is the subgraph of G induced by 3-vertices with 2-neighbors. A t -component of H is a component of order t .

Claim 3. *Each component of H is a path on at most 5 vertices.*

Suppose not. By construction, $\Delta(H) \leq 2$; since G has no 3-cycle, assume that some component H_1 of H induces a path or cycle $v_1 \dots v_k$ on 4 or more vertices. We consider a path first; the case of a cycle is similar and easier. Since Figure 2(a) is forbidden, every set of four successive 3-vertices on the path must contain a pair with a common 2-neighbor. Since G has no 3-cycles, no successive 3-vertices on the path have a common 2-neighbor. Similarly, since Figure 2(b) is forbidden, no vertices at distance two on the path have a common 2-neighbor. Thus, each path vertex v_i (for $i \in \{1, \dots, t-3\}$) must share a common 2-neighbor with v_{i+3} ; otherwise, we get the configuration in Figure 2(a) or Figure 2(b). This immediately gives that $t \leq 6$, since otherwise v_4 must share a common 2-neighbor with both v_1 and v_7 , a contradiction. If H_1 is a path on 6 vertices, then G is the Hajós join of P^* and a smaller graph J . Since G is 3-critical, Corollary 2 implies that J is also 3-critical. This contradicts our hypothesis. Thus, H_1 cannot be a path on 6 vertices.

To rule out a cycle, note that we can't pair up the 3-vertices so that v_i and v_{i+3} have a common 2-neighbor for each of the paths obtained by deleting a single cycle edge (since Figure 2(a) and Figure 2(b) need not be induced). If H_1 is a 4-cycle or 5-cycle, then any pairing gives a triangle or Figure 2(b); no pairing gives Figure 2(a). If H_1 is a 6-cycle, then we have only one possible pairing, but now the whole graph is P^* , which contradicts our hypothesis. Thus, H_1 must be a path on 5 or fewer vertices. \square

Claim 4. *No 3-vertex has two neighbors in the same component of H .*

Suppose that G contains such a 3-vertex v , and let H_1 be the component of H containing two of its neighbors. Claim 3 implies that H_1 is a path on at most 5 vertices; further, v must be adjacent to the endvertices of H_1 . If H_1 has 2 vertices, then G contains a 3-cycle, contradicting Claim 1. If H_1 has 3 vertices, then G contains Figure 2(c), contradicting

Claim 2. If H_1 has 4 vertices, then we can delete it and extend the coloring of the smaller graph using one of the two extensions shown in Figure 3, depending on which colors are available at the 2-vertices (if the color used on the edge incident to the 3-vertex is seen by both 2-vertices, we use the extension on the left; otherwise, the extension on the right). Finally, suppose that H_1 has five vertices. Now G is the Hajós join of P^* and a smaller graph J ; the copy of $P^* - e$ in G consists of H_1 , its adjacent 2-vertices, v , and v 's neighbor outside of H_1 . Corollary 2 implies that J is 3-critical. Since, J has lower average degree than G , it contradicts our choice of G as a minimal counterexample. \square




Figure 3: How to extend a coloring of $G \setminus H_1$ to G when H_1 has order 4 and a 3-vertex v has two neighbors in H_1 .

Claim 5. *The graph in Figure 4 cannot appear as a subgraph of G . Furthermore, the graph cannot appear as a subgraph even if one or more pairs of 2-vertices are identified. Thus, no 3-vertex has 3-neighbors in two distinct components of H , each of order at least 4.*

The final statement follows immediately from the first two. We defer the proofs of those two statements to Lemma 3 in Section 3. \square

Recall now our outline of the discharging proof in the introduction. To begin, each 2-vertex takes some charge from each 3-neighbor. Now 3-vertices with 2-neighbors need more charge and 3-vertices with no 2-neighbors have extra charge. Thus, we call a 3-vertex with a 2-neighbor *poor* and a 3-vertex with no 2-neighbor *rich*. Roughly, we give charge from rich vertices to poor vertices. A rich vertex is *type* (a, b, c) if it has 3-neighbors in components of




Figure 4: A configuration to forbid type $(*, 4^+, 4^+)$ vertices.

H of orders a, b, c . Typically, we choose a, b, c such that $a \leq b \leq c$. Analogous to vertices, a t^+ -component (resp. t^- -) has order at least (resp. at most) t . Claim 5 shows that vertices of type $(*, 4^+, 4^+)$ are forbidden. So Claims 3 and 4 imply that all rich vertices have type $(3^-, 3^-, 5^-)$.

Now we can finish the proof by discharging. Recall that each vertex v starts with charge $d(v)$. We apply the following three discharging rules, in succession.

(R1) Each 2-vertex takes charge $\frac{13}{37}$ from each neighbor.

(R2) Each rich vertex gives charge $\frac{t}{37}$ to each 3-neighbor in a t -component of H .

(R3) The 3-vertices in each component of H average their charge.

Now we verify that each vertex finishes with charge at least $2 + \frac{26}{37}$.

Each 2-vertex finishes with charge $2 + 2(\frac{13}{37}) = 2 + \frac{26}{37}$.

Consider a component H_1 of H . After (R1), each vertex has charge $3 - \frac{13}{37}$. By (R2) and (R3), each vertex gains $2(\frac{1}{37})$, so finishes with charge $3 - \frac{13}{37} + 2(\frac{1}{37}) = 2 + \frac{26}{37}$.

Since type $(*, 4^+, 4^+)$ vertices are forbidden, each rich vertex gives away charge at most $(3 + 3 + 5)(\frac{1}{37})$, so finishes with charge at least $3 - \frac{11}{37} = 2 + \frac{26}{37}$. \square

A common sentiment evoked by discharging proofs is that they're easy to verify, but hard to find. So to shed some light on this process of discovery, we conclude this section with a synopsis of how we found the proof of the Main Theorem.

We began not knowing what edge bound we could prove. That specific value came last. The general idea was to get all vertices as much charge as possible, say charge $2 + \alpha$, for some $\alpha \in (\frac{2}{3}, \frac{3}{4})$. Since each 2-vertex needs charge α , it takes charge $\frac{\alpha}{2}$ from each of its two 3-neighbors. Now 3-vertices with 2-neighbors have given away too much charge (thus, the name *poor*), so they need more charge from elsewhere. (If $\alpha = \frac{2}{3}$, then each 3-vertex gives away charge exactly $\frac{1}{3}$, so all vertices finish with charge $2 + \frac{2}{3}$, which proves Jakobsen's bound: $|E(G)| \geq \frac{4}{3}|V(G)|$. Recall that Vizing's Adjacency Lemma implies that each 3-vertex has at most one 2-neighbor.)

The poor vertices need extra charge, and they certainly won't get it from other poor vertices. Thus, we must show that each poor vertex has some nearby rich vertex. This motivates our definition of H . So far we have used no reducible configurations (only Vizing's Adjacency Lemma). By definition, each component of H is either a path or a cycle. Each endpoint of each path of H has a rich neighbor; crucially, rich vertices can share charge with the poor vertices in H (thus, the name *rich*). However, any cycle component of H has no such 3-neighbors to share charge with it. Thus, it is essential to show that H contains no cycles. Furthermore, it is helpful to show that each path in H is short, since the charge received by each path will be shared evenly among its vertices. To prove that H has the desired structure, we introduce the reducible configurations shown in Figure 2(a) and Figure 2(b). It is at this point that we first encounter P^* . If some component of H is a 6-cycle, then the entire graph G is P^* ; similarly if some component of H is a 5-path, with its endpoints having a common neighbor, then G is the Hajós join of P^* and a smaller 3-critical graph. We also simplify things a bit by showing that no path of H has endpoints with a common neighbor. The proof of this fact uses Figure 2(c).

To guarantee that each vertex of H finishes with charge at least $2 + \alpha$, we split the charge given to each component of H evenly among its vertices. Note that each path of H gets charge from two vertices. Since each vertex of H must reach final charge $2 + \alpha$, we have each t -component of H take from each rich neighbor a charge proportional to t . This leads to the definition of type (a, b, c) vertices. Of course, now we must bound the sum $a + b + c$ of a type (a, b, c) vertex that is not reducible. This leads to the reducible configuration in Figure 4, which shows that when $a \leq b \leq c$ we may assume that $b \leq 3$. Each component of H has order at most 5, so $a + b + c \leq 11$.

Following our framework above, each vertex must reach charge at least $2 + \alpha$, and each type (a, b, c) vertex gives charge $t\beta$ to each adjacent t -component of H . Thus, we choose α and β to maximize the minimum of the three expressions $2 + \alpha$, $3 - \frac{\alpha}{2} + 2\beta$, and $3 - 11\beta$. This maximum is attained when the three quantities are equal, at $\alpha = \frac{26}{37}$ and $\beta = \frac{1}{37}$. This proves the bound $2|E(G)| \geq (2 + \frac{26}{37})|V(G)|$.

3 Reducibility

In this section we prove Claim 2 and Claim 5 from Section 2, that certain subgraphs H , not necessarily induced, are forbidden in G .

In each case, by criticality we 3-color all but the edges of H (or some subgraph of it). If a color w is used on an edge incident to a vertex v_i , then v_i sees w . We want to show that we can always extend the partial coloring to all of G . We 3-color with the colors x , y , and z . Let $X = \{y, z\}$, $Y = \{x, z\}$, and $Z = \{x, y\}$. If a vertex v_i in H sees x , then the list of allowable colors for the uncolored edges incident to v_i is X ; similarly for colors y and z . If the subgraph H has t vertices each with one incident colored edge, then we write this as an ordered t -tuple, where each entry is X , Y , or Z . We call each possible t -tuple a *board*. By permuting color classes, we will assume that the first coordinate in every board is X , and the first coordinate different is Y . By Polya counting, the number of boards is $(3^{t-1} + 1)/2$.




Figure 5: A subgraph forbidden from appearing in 3-critical graph G .

Lemma 3. *The subgraph shown in Figure 5 (and in Figure 4) cannot appear in a 3-critical graph. Nor can it appear if we identify one or two vertex pairs in $\{v_6, v_7, v_{14}, v_{15}\}$.*

Proof. We first consider the case where no pairs of 2-vertices are identified. Note that the right and left sides of the figure are symmetric. By criticality, construct a partial 3-coloring




Figure 6: Extensions for part of Figure 5, based on the colors seen by v_2 , v_5 , v_6 , and v_7 .

of all of G except the edges incident to v_1 , v_3 , and v_4 . Since $t = 4$ (as defined in introduction to this section), we have $(3^3 + 1)/2 = 14$ boards. We begin by showing that for 12 of these 14 boards, we can extend the coloring to all of G .

If v_2 and v_5 see distinct colors, then Figure 6(a) shows how to extend the coloring unless the board is (X, X, Y, Y) : simply color greedily along the path of uncolored edges, starting at v_6 and ending at v_1 . Now suppose instead that v_2 and v_5 see the same color, x . If v_6 or v_7 sees x , then we can extend the coloring as in Figure 6(b): now color greedily along the path of uncolored edges, ending at v_7 . Further, if v_6 and v_7 see distinct colors, then we can color as in Figure 6(c). Thus, we conclude that we can extend the partial coloring to G unless the board is either (X, X, Y, Y) or (X, Y, Y, X) . Note that these two bad boards differ in the colors used on *two* edges, even up to all permutations of color classes. Thus, if at least one pendant edges is not yet colored, we can always find an extension of the partial coloring.

Now suppose that G contains a copy of Figure 5. By criticality, we get a 3-coloring of all of G except the edges with both endpoints in Figure 5. Our goal is to color the two remaining edges incident to v_8 so that both the left side and the right side can be colored from their resulting boards. As shown above, we must color v_5v_8 and v_8v_{10} so that neither the left or right board is (X, X, Y, Y) or (X, Y, Y, X) .

Given the colors incident to v_2 , v_6 , and v_7 , at most one choice of color for v_5v_8 gives a bad board for the left side. Similarly, at most one choice of color for v_8v_{10} is bad for the right side. We can color the edges as desired unless the color that is bad on v_5v_8 for the left side is the same as the color that is bad on v_8v_{10} for the right side, and that color, say x , is different from the color y seen by v_8 . So suppose this is true. Now perform an (x, y) -Kempe swap at v_8 . If this Kempe chain ends at neither the left nor right side, then we color v_5v_8 and v_8v_{10} arbitrarily. Now we can color each side. So suppose instead that the Kempe chain ends at the left side (by symmetry). Now we can color the left side, since v_5v_8 is uncolored. Afterward, the color for v_8v_{10} is determined, and we can color the right side. This completes the case where no pairs of 2-vertices are identified.

Now we consider the case where two vertex pairs in $\{v_6, v_7, v_{14}, v_{15}\}$ are identified. Since G has no 3-cycles, each of v_6, v_7 must be identified with one of v_{14}, v_{15} . By criticality, color all edges except those with both endpoints in Figure 5. We have only 3 incident colored edges, so 5 possible boards. Similar to above, we show how to extend the coloring to G ,

using Figure 6.

For board (X, X, X) , color v_5v_8 and v_4v_7 with y and color v_8v_{10} and $v_{11}v_{14}$ with z . (Perhaps $v_7 = v_{14}$, but this is okay.) Now we can extend the coloring to each side, as in Figure 6(a). A similar strategy works in every case except (X, Y, X) . We always color v_5v_8 and v_8v_{10} so that their colors differ from those seen by v_2 and v_{13} , respectively. Next, we color v_4v_7 and $v_{11}v_{14}$ to match v_5v_8 and v_8v_{10} , respectively. Finally, we can color each side as in Figure 6(a). So consider case (X, Y, X) . Color v_5v_8 with x and v_8v_{10} with z . Now color v_3v_4 with x and $v_{14}v_i$ with y , where $i \in \{3, 4\}$. Since v_{10} and v_{14} see different colors, we extend the right side as in Figure 6(a). We extend the left side as in Figure 6(b). This completes the case of two pairs of identified vertices.

Now suppose that one vertex pair in $\{v_6, v_7, v_{14}, v_{15}\}$ is identified; we consider three cases. The identified pair is either (v_6, v_{15}) , (v_7, v_{14}) , or (v_6, v_{14}) ; we call these cases “outside”, “inside”, and “mixed”. In each case, five vertices see colors, but we initially consider only the colors seen by v_2 , v_8 , and v_{13} . Thus, for example, we write the board (X, Y, Z) to signify that v_2 sees x , v_8 sees y , and v_{13} sees z .

First consider outside. Suppose we have board (X, X, Y) . Color v_3v_6 with x , color v_5v_8 and $v_{12}v_{15}$ with y , and color v_8v_{10} with z . We can extend the coloring on each side as in Figure 6(a). A similar strategy works for boards (X, Y, Y) and (X, Y, Z) . Consider instead (X, Y, X) . Now color v_5v_8 and $v_{12}v_{15}$ with x and color v_8v_{10} with z . We can color the right as in Figure 6(a) and the left as in Figure 6(b). Finally, consider (X, X, X) . If v_7 sees x , then color $v_{12}v_{15}$ with x , color v_5v_8 with y , and color v_8v_{10} with z . Now color both sides as in Figure 6(a). Otherwise, by symmetry v_7 sees y . Now color $v_{12}v_{15}$ with x , v_8v_{10} with y , and v_5v_8 with z . We can again color both sides as in Figure 6(a).

Now consider inside. This case is similar to above. Consider a board *other* than (X, Y, X) . Color v_5v_8 and v_8v_{10} to differ from the colors seen by v_2 and v_{13} , respectively. Now color v_4v_7 and $v_{11}v_{14}$ to match v_5v_8 and v_8v_{10} , respectively. Finally, color each side as in Figure 6(a). So consider (X, Y, X) . If v_6 sees x , then color v_5v_8 with x and v_8v_{10} with z . Now we can color the right first, then color the left, since the left board is $(X, X, *, X)$. So v_6 does not see x . Now color v_5v_8 with z and v_8v_{10} with x . Color the right first, then color left as in Figure 6(a), since v_2 and v_6 see distinct colors.

Finally, consider mixed. Recall that v_6 and v_{14} are identified. Consider a board other than (X, X, X) and (X, Y, Y) . If v_{13} and v_{15} see distinct colors, then color v_8v_{10} with a color not seen by v_{13} . Now color the left side, then extend to the right as in Figure 6(a). Otherwise v_{13} and v_{15} see the same color, so use that color on v_8v_{10} . Now color the left, then extend to the right, as in Figure 6(b). Instead, consider (X, X, X) . Color v_3v_6 with x , color v_5v_8 with y , and both v_8v_{10} and $v_{11}v_{14}$ with z . Now extend both sides as in Figure 6(a). Finally, consider (X, Y, Y) . If v_{15} sees a color other than y , then color v_8v_{10} to avoid the color seen by v_{15} . Now color the left, followed by the right, as in Figure 6(a). Similarly, if v_7 sees x , then color v_5v_8 with x and color the right, followed by the left. Likewise, if v_7 sees y , then color v_5v_8 with z and color the right, followed by the left. Thus, we conclude that v_7 sees z and v_{15} sees y . Now perform an (x, y) -Kempe swap at v_8 . The resulting board will be one of the cases above. \square




Figure 7: This subgraph cannot appear in a 3-critical graph.

Lemma 4. *The subgraph in Figure 7 (and Figure 2(a)) cannot appear in a 3-critical graph.*

Proof. To describe the boards, we use an ordered 6-tuple, where coordinate i is the list of colors missed by v_i . By Polya counting, the number of boards is $(3^5 + 1)/2 = 122$. The number of these of type $(X, *, *, *, *, X)$ is $(3^4 + 1)/2 = 41$. In the remaining $122 - 41 = 81$ boards, the first and last coordinates differ. By symmetry of color classes, we denote them as type $(X, *, *, *, *, Y)$. Similar to the previous proof, we seek to color edge v_7v_8 , so that the resulting boards for the left and right side are both colorable. (Some example colorings are shown in Figure 8.) However, some boards are not immediately colorable. Thus, we sometimes perform one or more Kempe swaps on a board before the result is colorable. If, from a given board, we can always perform some sequence of Kempe swaps to reach a colorable board, then we *win* on that board. Suppose that vertices v_i and v_j each see exactly one of colors x and y . If v_i and v_j lie in the same (x, y) -Kempe chain, then they are (x, y) -paired. If a vertex v_i is not (x, y) -paired with any vertex v_j , then it is (x, y) -unpaired. The definitions for color pairs (x, z) and (y, z) are analogous. To prove that we can win on all boards, we consider separately type $(X, *, *, *, *, X)$ and type $(X, *, *, *, *, Y)$.

Case 1: The board has type $(X, *, *, *, *, X)$. We group boards based on how many leading X s they have. This gives five possibilities: (X, X, X, X, X, X) ; (X, X, X, X, Y, X) ; $(X, X, X, Y, *, X)$; $(X, X, Y, *, *, X)$; and $(X, Y, *, *, *, X)$. The first two types can be colored using two copies of Figure 8(a). Two instances of the third type can be colored again using two copies of Figure 8(a). The remaining instance of this type is not colorable: (X, X, X, Y, Z, X) . Among type $(X, X, Y, *, *, X)$, seven are colorable using Figure 8(a) and either Figure 8(a) or Figure 8(b). This leaves (X, X, Y, Y, Z, X) and (X, X, Y, Z, Y, X) . The first is colorable using Figure 8(e) and a reflected copy of Figure 8(e). The second using Figure 8(f) and a reflected copy of Figure 8(f).

Now we consider type $(X, Y, *, *, *, X)$. Of these 27 boards, seven of type $(X, Y, X, *, *, X)$ are colorable using Figure 8(a) and either Figure 8(a) or Figure 8(b). Similarly, seven of type $(X, Y, Y, *, *, X)$. Also, (X, Y, Y, Z, Y, X) is colorable using Figure 8(f) and a reflected copy of Figure 8(f). Four of type $(X, Y, Z, *, *, X)$ are colorable using Figure 8(f) and a reflected copy of Figure 8(f). Combining these eight uncolorable boards with the one in the previous paragraph, we conclude that type $(X, *, *, *, *, X)$ is colorable unless it is one of the following nine boards: (X, X, X, Y, Z, X) ; (X, Y, X, Y, Z, X) ; (X, Y, X, Z, Y, X) ; (X, Y, Y, Y, Z, X) ; (X, Y, Z, X, X, X) ; (X, Y, Z, X, Y, X) ; (X, Y, Z, X, Z, X) ; (X, Y, Z, Y, Z, X) ; (X, Y, Z, Z, Z, X) .

Consider the eighth of these: (X, Y, Z, Y, Z, X) . Now we (y, z) -swap at v_2 (regardless of its




Figure 8: Seven valid colorings for half of the configuration. Three of the four pendant edges in each figure represent pre-colored edges, incident to the configuration. The fourth pendant edge will be v_7v_8 . Each valid coloring of the configuration (for some choice of board) consists of two figures that share a color on the edge where they overlap. For example, the board (X, Y, Z, X, Y, Y) can be colored using Figures 8(f) on the left and 8(g) on the right.

(y, z) -pair). After permuting color classes, the result is (X, Y, Z, Z, Y, X) or (X, Y, Y, Y, Y, X) or (X, Y, Y, Z, Z, X) . Each possibility is colorable, as shown in the previous paragraph. The remaining eight uncolorable boards split into forwards/backwards pairs (which are equivalent, since the configuration is symmetric). These are first and fifth; second and seventh; third and sixth; and fourth and ninth. We note also that second and third are equivalent; if we can win on the second, then we can win on the third, and vice versa. Consider the third: (X, Y, X, Z, Y, X) . Now we (y, z) -swap at v_4 . If (v_2, v_4) is a (y, z) -pair or if v_4 is (y, z) -unpaired, then we reach a colorable board. Otherwise, (v_4, v_5) is a (y, z) -pair, so we reach (X, Y, X, Y, Z, X) , the second. A similar argument shows that we can reach the third from the second. Thus it suffices to show that we can win on boards (X, X, X, Y, Z, X) , (X, Y, X, Y, Z, X) , and (X, Y, Y, Y, Z, X) .

We first consider board (X, Y, X, Y, Z, X) . Now we (x, y) -swap at v_4 (regardless of its (x, y) -pair). If (v_4, v_2) or (v_4, v_3) is an (x, y) -pair, then the resulting board is colorable, as shown above; similarly if v_4 is (x, y) -unpaired. If, instead, (v_1, v_4) is an (x, y) -pair, then the result is (Y, Y, X, X, Z, X) , which is (X, X, Y, Y, Z, Y) . This can be colored using B1 and A4. Suppose instead that (v_4, v_6) is a (x, y) -pair. Now (x, y) -swapping (v_4, v_6) yields (X, Y, X, X, Z, Y) , which is colorable Figure 8(a) and Figure 8(e). Thus, we win on (X, Y, X, Y, Z, X) .

We next consider board (X, X, X, Y, Z, X) . We (x, z) -swap v_3 if either v_3 is (x, z) -unpaired or (v_2, v_3) or (v_3, v_5) is an (x, z) -pair. In each case, the result is colorable, as shown

above. So assume that either (v_3, v_6) or (v_3, v_1) is an (x, z) -pair. In the first case, we (x, z) -swap (v_3, v_6) to get (X, X, Y, Z, Y, Y) (after permuting color classes). This board is colorable using Figure 8(e) and Figure 8(d). So assume that (v_3, v_1) is an (x, z) -pair. Now consider the (x, z) -pair of v_5 . If v_5 is (x, z) -unpaired, then we (x, z) -swap v_5 to get (X, X, X, Y, X, X) , which is colorable using two copies of Figure 8(a). If (v_5, v_2) is an (x, z) -pair, then we (x, z) -swap (v_5, v_2) , to get (X, Y, X, Z, X, X) (up to symmetry), which is colorable using Figure 8(a) and Figure 8(b). So (v_5, v_6) must be an (x, z) -pair, which means that v_2 is (x, z) -unpaired. So we (x, z) -swap v_2 to get (X, Y, X, Z, Y, X) (up to symmetry). As shown above, this board is equivalent to (X, Y, X, Y, Z, X) , which we showed in the previous paragraph is a win.

Finally, we consider board (X, Y, Y, Y, Z, X) . Now we (x, y) -swap v_4 , as long as (v_4, v_1) is not an (x, y) -pair. If v_4 is (x, y) -unpaired, this gives (X, Y, Y, X, Z, X) ; if (v_2, v_4) is a (x, y) -pair, this gives (X, X, Y, X, Z, X) ; if (v_3, v_4) is a (x, y) -pair, this gives (X, Y, X, X, Z, X) . Each of these boards is colorable using Figure 8(a) and Figure 8(b). If (v_4, v_6) is an (x, y) -pair, then we (x, y) -swap (v_4, v_6) to get (X, Y, Y, X, Z, Y) , which is colorable using Figure 8(a) and Figure 8(e). So assume that (v_1, v_4) is an (x, y) -pair. If (v_2, v_3) is also an (x, y) -pair, then we (x, y) -swap there to get (X, X, X, Y, Z, X) , on which we win, by the previous paragraph. If v_2 has no (x, y) -pair, then we (x, y) -swap at v_2 to get (X, X, Y, Y, Z, X) , which is colorable using Figure 8(e) and a reflected copy of Figure 8(e). Thus (v_2, v_6) is an (x, y) -pair. So v_3 is (x, y) unpaired; now (x, y) -swapping at v_3 gives (X, Y, X, Y, Z, X) , on which we win, as shown above. Hence, we win on all boards of type $(X, *, *, *, *, *, X)$.

Case 2: The board has type $(X, *, *, *, *, Y)$. We first show that we can win all boards of type $(X, X, Y, *, *, Y)$. In fact, all such boards are colorable. Seven of the nine colorings use Figure 8(e) with either Figure 8(c) or Figure 8(d). One of the remaining colorings uses Figure 8(a) and Figure 8(e). the other uses Figure 8(f) and Figure 8(g). An analogous argument (with Figure 8(b) in place of Figure 8(a)) shows that we win all boards of type $(X, X, Z, *, *, Y)$.

Now consider boards of type $(X, X, X, *, *, Y)$. Four of these are colorable using Figure 8(a) and Figure 8(e). So we consider five boards: (X, X, X, X, X, Y) ; (X, X, X, Y, X, Y) ; (X, X, X, Y, Y, Y) ; (X, X, X, Y, Z, Y) ; (X, X, X, Z, X, Y) . In each case, (v_1, v_6) must be an (x, y) -pair; otherwise, we (x, y) -swap at v_6 and win by Case 1. Similarly, (v_2, v_3) must be an (x, y) -pair; otherwise, we (x, y) -swap at v_3 , and reduce to a board of type $(X, X, Y, *, *, Y)$, which we win as above. In the fourth and fifth cases, we (x, y) -swap v_4 and v_5 , respectively; in each case, the resulting board can be colored using Figure 8(a) and Figure 8(e). Hence, we are in one of the first three cases, so (v_4, v_5) is an (x, y) -pair. In the second case, we (x, y) -swap (v_4, v_5) ; this yields (X, X, X, X, Y, Y) , which is colorable using Figure 8(a) and Figure 8(e). In the first and third cases, we (x, y) -swap (v_2, v_3) and possibly (v_4, v_5) , in each case yielding (X, Y, Y, X, X, Y) . This is colorable using Figure 8(f) and Figure 8(g). Thus, we win all boards of type $(X, X, *, *, *, Y)$.

Now consider a board of type $(X, Y, *, *, *, Y)$. We can assume that (v_1, v_6) is an (x, y) -pair, for otherwise we (x, y) -swap at v_6 and can win by Case 1. However, now we (x, y) -swap v_2 . This yields a board of type $(X, X, *, *, *, Y)$, which we can win, as shown above. Thus, we win all boards of type $(X, Y, *, *, *, Y)$.

Now we consider boards of type $(X, Z, *, *, *, Y)$. First consider type $(X, Z, Z, *, *, Y)$. Seven of these nine boards are colorable using Figure 8(e) with either Figure 8(c) or Fig-

ure 8(d). This leaves (X, Z, Z, X, Z, Y) and (X, Z, Z, Z, X, Y) . The first is colorable using Figure 8(b) and Figure 8(e), so consider the second. Now (v_1, v_2) must be an (x, z) -pair; otherwise, we (x, z) -swap at v_2 and reduce to type $(X, X, *, *, *, Y)$. If either v_4 or v_5 is (x, z) -unpaired, then we (x, z) -swap there and reduce to a colorable board. Hence, (v_4, v_5) is a (x, z) -pair. Now, again, we (x, z) -swap (v_4, v_5) and reduce to (X, Z, Z, X, Z, Y) , which is colorable. Thus, we win all boards of type $(X, Z, Z, *, *, Y)$.

Now consider type $(X, Z, Y, *, *, Y)$. Note that (v_2, v_6) must be a (y, z) -pair; otherwise, we (y, z) -swap v_2 and reduce to the case $(X, Y, *, *, *, Y)$. However, now we (y, z) -swap v_3 , and reduce to the case $(X, Z, Z, *, *, Y)$.

Finally, consider type $(X, Z, X, *, *, Y)$. Recall that (v_1, v_6) must be an (x, y) -pair; otherwise we (x, y) -swap at v_6 and win by Case 1. Now we (x, y) -swap at v_3 , and win by the previous paragraph. So we win all boards of type $(X, Z, *, *, *, Y)$. Thus, we win all boards of type $(X, *, *, *, *, Y)$, which completes the proof of both the claim and the lemma. \square

4 An Improved Bound using a Computer




Figure 9: Extra reducible configurations.

The reducible configurations in this paper were originally found by computer. The computer uses an abstract definition capturing the notion of “colorable after performing some Kempe swaps”, which frees it from considering an embedding in an ambient critical graph. This is called *fixability* and extends the idea in [6] from stars to arbitrary graphs. The computer is able to prove many reducibility results for which we have yet to find short proofs. Here we show how to use some of these reducible configurations to further improve the bound on the average degree of 3-critical graphs. We give a larger survey of what can be proved with these computer results in [1].

Lemma 5. *The configurations in Figure 9 cannot be subgraphs of a 3-critical graph. In particular, all rich vertices are of type $(2^-, 2^-, 5^-)$.*

Proof. Suppose that v is a rich vertex of type $(*, 3^+, 3^+)$. If these two adjacent components of H have no common 2-neighbors, then G has a copy of Figure 9(i). Otherwise, the components share one or more common 2-neighbors, and G contains one of Figure 9(a)–(h). The computer is able to generate proofs in L^AT_EX, but at about 100 pages this one is not a fun read: <https://dl.dropboxusercontent.com/u/8609833/Papers/big%20tree.pdf> \square

Theorem 6. *If a graph G with $\Delta = 3$ is critical, then either $2|E(G)| \geq (2 + \frac{22}{31})|V(G)|$ or else either G is P^* or J_1 (Woodall’s first example).*

Proof. By Lemma 5, we have $a + b + c \leq 9$ for every rich vertex of type (a, b, c) . So, we obtain our desired bound when the three quantities $2 + \alpha$, $3 - \frac{\alpha}{2} + 2\beta$ and $3 - 9\beta$ are equal. This happens when $\alpha = \frac{22}{31}$ and $\beta = \frac{1}{31}$, which yields $2|E(G)| \geq (2 + \frac{22}{31})|V(G)|$. \square

References

- [1] Daniel W. Cranston and Landon Rabern, *Edge-coloring via Fixable Subgraphs*, In preparation (2015).
- [2] S. Fiorini and Robin James Wilson, *Edge-colourings of graphs*, Pitman, London; distributed by Fearon-Pitman Publishers, Inc., Belmont, Calif., 1977, Research Notes in Mathematics, No. 16.
- [3] Stanley Fiorini and Robin J. Wilson, *On the chromatic index of a graph. II*, Combinatorics (Proc. British Combinatorial Conf., Univ. Coll. Wales, Aberystwyth, 1973), Cambridge Univ. Press, London, 1974, pp. 37–51. London Math. Soc. Lecture Note Ser., No. 13.
- [4] Ivan Taftberg Jakobsen, *Some remarks on the chromatic index of a graph*, Arch. Math. (Basel) **24** (1973), 440–448.
- [5] ———, *On critical graphs with chromatic index 4*, Discrete Math. **9** (1974), 265–276.
- [6] Landon Rabern, *A game generalizing Hall’s theorem*, Discrete Math. **320** (2014), 87–91.
- [7] Michael Stiebitz, Diego Scheide, Bjarne Toft, and Lene M. Favrholdt, *Graph edge coloring*, Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2012, Vizing’s theorem and Goldberg’s conjecture, With a preface by Stiebitz and Toft.
- [8] V. G. Vizing, *The chromatic class of a multigraph*, Kibernetika (Kiev) **1965** (1965), no. 3, 29–39.
- [9] ———, *Some unsolved problems in graph theory*, Uspehi Mat. Nauk **23** (1968), no. 6 (144), 117–134.
- [10] Douglas R. Woodall, *The average degree of an edge-chromatic critical graph*, Discrete Math. **308** (2008), no. 5-6, 803–819.